arXiv:1510.07357v1 [cs.DC] 26 Oct 2015

Distributed Communication in Bare-Bones Wireless Networks

Bogdan S. Chlebus * Dariusz R. Kowalski † Shailesh Vaya ‡

Abstract

We consider wireless networks in which the effects of interference are determined by the SINR model. We address the question of structuring distributed communication when stations have very limited individual capabilities. In particular, nodes do not know their geographic coordinates, neighborhoods or even the size n of the network, nor can they sense collisions. Each node is equipped only with its unique name from a range $\{1, \ldots, N\}$. We study the following three settings and distributed algorithms for communication problems in each of them. In the uncoordinated-start case, when one node starts an execution and other nodes are awoken by receiving messages from already awoken nodes, we present a randomized broadcast algorithm which wakes up all the nodes in $\mathcal{O}(n \log^2 N)$ rounds with high probability. In the synchronized-start case, when all the nodes simultaneously start an execution, we give a randomized algorithm that computes a backbone of the network in $\mathcal{O}(\Delta \log^7 N)$ rounds with high probability. Finally, in the partly-coordinated-start case, when a number of nodes start an execution together and other nodes are awoken by receiving messages from the already awoken nodes, we develop an algorithm that creates a backbone network in time $\mathcal{O}(n \log^2 N + \Delta \log^7 N)$ with high probability.

Key words: Wireless network, signal-to-interference-plus-noise ratio, broadcast, knowledge, backbone structure

*Department of Computer Science and Engineering, University of Colorado Denver, Denver, Colorado, U.S.A. Supported by the NSF Grant 1016847.

†Department of Computer Science, University of Liverpool, Liverpool, United Kingdom.

‡Xerox Research Centre India, Bangalore, India.

1 Introduction

We consider ad-hoc wireless networks in which the effects of interference are determined by the signal-to-interference-plus-noise ratio (SINR) model. The extent to which such networks can support distributed communication depends on the capabilities of nodes (such as collision detection) and information that can be used in codes of algorithms (such as the geographic coordinates of nodes). This paper explores how restricted nodes can and still participate in efficient distributed communication in wireless networks, subject to certain technical specifications of the wireless model.

Nodes of a network are placed in a two-dimensional Euclidean space. What we call the communication graph of a network has the network's nodes as graph vertices and a pair of nodes (u, w) make an edge when w can successfully receive a message transmitted by u when no other node is transmitting simultaneously. We denote the maximum degree of a node in the communication graph by Δ. Nodes cannot detect collisions produced by interfering transmissions. We work with the SINR model that combines the weak connectivity assumptions formulated by Daum et al. [4] and the weak sensitivity (weak devices) assumptions formulated by Jurdziński et al. [17], see also [12].

Algorithms are restricted in what information can be used in their codes. We call *bare-bones* the specific set of such restrictions used in this paper, which can be summarized as follows. Each among the n nodes in the network has a unique name in the range $\{1, \ldots, N\}$, for some positive integer $N \geq n$. Each node knows its name and the number N. The nodes do not know their neighbors in the communication graph but they know the parameter Δ.

We consider distributed algorithms to synchronize and organize a wireless network. Algorithms' design and performance depend on how a communication operation is initiated. Executing an algorithm from an uncoordinated start means that either just a single node is awoken at the beginning of an execution, while all other nodes stay passive until they hear a message, or a group of nodes start an execution together while other nodes need to be awoken by hearing messages. When all the nodes begin an execution together in the same round, then such an execution is said to be performed from a synchronized start.

We consider three tasks to be carried out by distributed algorithms. One task is to prepare a synchronized start from an *uncoordinated start*, when just one node is awoken spontaneously. Another task is to prepare a synchronized start from a *partly coordinated start*, when any number of nodes are awoken simultaneously while other nodes need to hear a message to become awoken. These tasks can be accomplished by a wake-up broadcast which also coordinates the clocks of nodes. The third task is to build a backbone of a network, provided a *synchronized start* is already reached. Backbone nodes induce a subgraph of communication graph of constant degree with asymptotically the same diameter as the underlying communication network and such that every node outside the backbone is connected to some node in the backbone.

A summary of the results. We develop a randomized algorithm which performs a broadcast from an uncoordinated start in $\mathcal{O}(n \log^2 N)$ expected time. Its execution results in waking up every node and synchronizing their clocks so that they can simultaneously start a computation, for instance, to find a backbone. The time complexity of this synchronization algorithm, from an uncoordinated start, is only $\mathcal{O}(\log n)$ away from the known lower bound given by Jurdziński et al. [17].

We present a randomized algorithm that computes a backbone from a synchronized start in $\mathcal{O}(\Delta \log^7 N)$ rounds with high probability. Finally, we develop an algorithm that creates a backbone network, for *any* non-empty set of initially activated nodes, in $\mathcal{O}(n \log^2 N + \Delta \log^7 N)$ rounds with high probability. In all presented algorithms, a node uses $\mathcal{O}(\log^3 N)$ random bits in an execution.

There are known lower bounds $\Omega(n \log N)$ and $\Omega(D\Delta)$ on time to broadcast from an uncoordinated start, even in restricted classes of networks, where D is the diameter of the network, given by Jurdziński et al. [17]. These lower bounds also hold for randomized algorithms and even for a setting when nodes may know their coordinates, so they hold for a more restricted setting considered in this work. Observe that building a backbone from an uncoordinated start is at least as hard as coordinating all nodes so that they can start an execution simultaneously. For $D = \Delta = \Theta(\sqrt{n})$, both backbone and broadcast could be performed in $\mathcal{O}(\Delta \text{ polylog } N) = \mathcal{O}(\sqrt{n} \text{ polylog } N)$ rounds, while for the uncoordinated start they may need as much as $\Omega(D\Delta) = \Omega(n)$ rounds. The performance of our algorithm for building a backbone from a synchronized start implies that the task of preprocessing a network to synchronize its nodes, so that they can start an execution simultaneously, improves time-performance of some nontrivial communication tasks, like building a virtual backbone network.

Our work could be seen as an extension of the results of Jurdziński and Kowalski [11] and Jurdziński et al. [17], whose solutions rely in an essential manner on the knowledge of coordinates, to a setting when such knowledge is not available. These are the first algorithms with performance close to optimal, in view of lower bounds in a less restricted setting [17] and which apply in our setting as well, in the wireless-network model of weak sensitivity and weak connectivity. To compensate for the lack of geographic node coordinates, we introduce new techniques of collision avoidance based on strongly selective families with specifically chosen parameters. We combine them with a number of graph-based and geometry-based algorithmic techniques to synchronize and build a backbone structure in wireless networks.

Previous work. Work on communication in SINR wireless networks is notorious for the variety of assumptions made by authors to specify settings that make efficiency of communication algorithms developed for such variants sometimes impossible to compare directly. We categorize communication problems following a methodology of SINR wireless networks popularized by Jurdziński and Kowalski [12]. It rests on two independent assumptions. One is about weak versus strong sensitivity, which is also known as weak versus strong devises, according to the terminology used by Jurdziński et al. [17]. The other is weak versus strong links in reachability graphs, according to the terminology introduced by Daum et al. [4]. These assumptions together determine four different settings, which we explain in detail in Section 2. Combined with these, there are also assumptions about which parameters of networks are known by nodes, so that they can be used in codes of algorithms.

A setting closest to the one of this work, is the *weak-sensitivity and weak-connectivity* model considered by Jurdziński et al. [17]. They proved lower bounds $\Omega(n \log N)$ and $\Omega(D\Delta)$ on time to broadcast from an uncoordinated start. These lower bounds hold for randomized algorithms and for a setting when nodes know their geographic information. They also developed a deterministic algorithm that accomplishes broadcast in time $\mathcal{O}(\min\{D\Delta \log^2 N, n \log N\})$. For the same setting, Jurdziński and Kowalski [11] gave a deterministic distributed algorithms building a backbone structure from a synchronized start in $\mathcal{O}(\Delta \text{ polylog } n)$ rounds.

For the model of *weak sensitivity and strong connectivity*, Yu et al [23] gave broadcast algorithms operating in times $\mathcal{O}(D + \log^2 n)$ and $\mathcal{O}(D \log^2 n)$ with high probability, where a specific bound depends on how broadcast is initiated. These algorithms additionally resort to a power-control mechanism. This approach applies to broadcasting multiple messages and was generalized by Yu et al. [24] to a scenario when nodes are activated in arbitrary rounds. The model of wireless networks in [23, 24] uses additional assumptions on the network parameters on which fast algorithms rely; one could expect that in such environments even more involved communication and computation tasks have solutions scalable in D.

For the model of *strong sensitivity and weak connectivity*, Daum et al. [4] gave a broadcast algorithm operating in time $\mathcal{O}(n \log^2 n)$. They showed a lower bound $\Omega(n)$, which holds in networks of diameter 2.

Finally, for the model of *strong sensitivity and strong connectivity*, Jurdziński et al. [13] gave a broadcast algorithm that works in time $\mathcal{O}(D \log n + \log^2 n)$ with high probability, which relies on nodes knowing the geographic coordinates. Jurdziński et al. [14] gave another algorithm that works in time $\mathcal{O}(D \log^2 n)$ with high probability. The latter algorithm does not rely on the knowledge of geographic coordinates, improving in most ranges of model parameters the performance of algorithms for this model given by Daum et al. [4]. The solution in [14] was generalized to the wake-up problem with non-synchronized start-ups by Jurdziński et al. [15]. For this model of strong sensitivity and strong connectivity, Jurdziński et al. [16] studied deterministic solutions for the single broadcast problem when nodes know their own spacial coordinates and those of their neighbors. Their deterministic algorithm from an uncoordinated start has $\mathcal{O}(D \log^2 n)$-time performance and another deterministic algorithm from a synchronized start operates in time $\mathcal{O}(D + \log^2 n)$. It remains open to develop deterministic algorithms for this model of comparable time performance that do not rely on knowledge of spatial coordinates.

Scheideler et al. [21] gave an algorithm to find a dominating set that stabilizes in time $\mathcal{O}(\log n)$. That algorithm relies on manipulating thresholds in physical carrier sensing. Constructions and applications of virtual backbones and dominating sets in wireless networks were surveyed by Yu et al. [26].

Related work. Local broadcasting, in which the requirement is to inform only the immediate neighbors in the corresponding communication graph, was studied by Goussevskaia et al. [8], Kesselheim and Vöcking [18], Yu et al. [25], Halldórsson and Mitra [10] and Fuchs and Wagner [6]. Centralized algorithms for the SINR model were surveyed by Goussevskaia et al. [9]. Broadcasting of multiple messages was considered by Reddy et al. [20] and Yu et al. [23, 24].

There has been a great variety of trends in work on models of signal strength and geometric decay in representations of wireless networks. Some authors have sought to specialize SINR-like settings to the Euclidean space to leverage its specific topological properties, like Avin et al. [1] through their work on SINR diagrams. Others have sought to abstract from geographic coordinates but maintain a metric space of bounded growth that defines distances, like that of Daum et al. [4] and Jurdziński et al. [14]. Still others, like Bodlaender and Halldórsson [2], attempted to abstract the geometry even further and consider models that relax the properties of distance as represented by metric spaces.

2 Technical Preliminaries

Executions of algorithms are assumed to be synchronous, in that they are structured as sequences of rounds of equal duration. The size of messages transmitted by nodes are scaled to the duration of rounds, so that one message can be transmitted and received in a round. A node acts either as a transmitter or a receiver in each round. Numbers of rounds, as determined by an external global clock, are not used.

There are n nodes in a network. Nodes have unique names assigned to them. Each node's name is an integer in the range $[1, N]$, for some positive integer N. We treat the numbers n and N as independent parameters in formulas, subject to the only assumption that $N \geq n$.

A broadcast algorithm disseminates some contents, to which we refer as a *rumor*, throughout the network. Messages are structured to consist of at most one rumor and a sequence of control bits that help coordinate concurrent actions of nodes. This model is often referred to in the literature as of *separate* or *bounded* messages. The number of control bits per message in our algorithms is $\mathcal{O}(\log N)$.

The model of wireless networks. Networks are embedded in a two-dimensional Euclidean space which determines the distance $\mathrm{dist}(x, y)$ for any pair of points.

The SINR interference model involves the following parameters: a *path-loss parameter* $\alpha > 2$, an *ambient-noise parameter* $\mathcal{N} > 0$, a *threshold* $\beta \geq 1$, a *sensitivity parameter* $\varepsilon_s > 0$, and a *connectivity parameter* ε_c such that $0 \leq \varepsilon_s \leq \varepsilon_c < 1$. In specific analyses, we assume that $\beta = 1$, but that is for the sake of simplicity and β can be scaled to any value $\beta \geq 1$.

The *transmission power* of a node v is a positive real number denoted by P_v. Let us suppose that a set of nodes $\mathcal{T}$ transmit in a round. For any two nodes v and u, where $v \in \mathcal{T}$ and $u \notin \mathcal{T}$, the *signal-to-interference-plus-noise ratio* in this round is denoted as $\mathrm{SINR}(v, u, \mathcal{T})$ and is defined as follows:

$$\mathrm{SINR}(v, u, \mathcal{T}) = \frac{P_v \, \mathrm{dist}(v, u)^{-\alpha}}{\mathcal{N} + \sum_{w \in \mathcal{T} \setminus \{v\}} P_w \, \mathrm{dist}(w, u)^{-\alpha}} \ . \tag{1}$$

We consider only *uniform* networks in which all transmission powers are equal to some number P.

Next we explain the categorization of models with respect of sensitivity. Again, let a set of nodes $\mathcal{T}$ transmit in a round and let nodes v and u be such that $v \in \mathcal{T}$ and $u \notin \mathcal{T}$. In the model of *strong sensitivity*, the node u receives a message from v successfully in this round when the inequality $\mathrm{SINR}(v, u, \mathcal{T}) \geq \beta$ holds. In such a situation, the largest range from which a message can be received with only one node transmitting equals $r = (P/(\mathcal{N}\beta))^{1/\alpha}$. This number is the result of maximizing $\mathrm{dist}(v, u)$ in Equation (1) in the environment most conducive to transmitting, as determined by the given noise $\mathcal{N}$ and the absence of interference power $\sum_{w \in \mathcal{T} \setminus \{v\}} P_w \, \mathrm{dist}(w, u)^{-\alpha}$ of other possible transmissions in the round, that is, when

$$\sum_{w \in \mathcal{T} \setminus \{v\}} P_w \, \mathrm{dist}(w, u)^{-\alpha} = 0 \ .$$

In the model of *weak sensitivity*, the node u successfully receives a message from v in this round when the following two inequalities hold: $\mathrm{SINR}(v, u, \mathcal{T}) \geq \beta$ and $\mathrm{dist}(v, u) \leq (1 - \varepsilon_s)r$.

The categorization of models into these two classes, of weak and strong sensitivities (devices), was introduced by Jurdziński et al. [17]. The motivation for weak sensitivity is that it is often too

costly for wireless devices to perform signal acquisition continuously, see [7]. An alternative is for the nodes to wait for an energy spike, as represented by the condition $\mathrm{SINR}(v, u, \mathcal{T}) \geq \beta$. Once the nodes experience it, they start sampling and correlating to synchronize and acquire a potential packet preamble [22]. After that they can detect signals as determined by the formulas of weak sensitivity. In this paper we use the model of weak sensitivity.

The *communication graph* is defined as follows: all the network's nodes are its vertices, and for any two nodes u and v, they are connected by an edge when the inequality $\mathrm{dist}(u, v) \leq (1 - \varepsilon_c)r$ holds. The maximum degree of a node in such a graph is denoted by Δ. When $\varepsilon_c = \varepsilon_s$ then the model is of *weak connectivity*, and if $\varepsilon_c > \varepsilon_s$ then the model is of *strong connectivity*. The categorization of wireless networks with respect to weak and strong connectivity was introduced by Daum et al. [4].

In this paper we use the model of weak connectivity. There are two reasons that motivate focusing on this model. First, the strong-connectivity model is not well scalable with the diameter of the communication graph, as shown, for example, in [4]. Second, the topology of the communication graph in the weak-connectivity model is less affected by the physical parameters of wireless communication that are not reflected in the SINR methodology.

What we call *bare-bones* wireless communication is a combination of the following restrictions on algorithms: the messages are separate and each node knows only its name and the numbers N and Δ. The motivation for studying such algorithms is that they operate under severe restrictions imposed by the size of messages and the available information and so are of wide applicability. Algorithms with similar restrictions were considered by Jurdziń ski et al. [14] in the case of strong-sensitivity and strong-connectivity of the wireless model, while we consider bare-bones algorithms for the combination of weak sensitivity and weak connectivity of wireless networks. The algorithms in [11, 17], which are also for weak-connectivity and weak-sensitivity of the wireless communication, are designed with the assumption that the nodes know their spatial coordinates; in this work we develop efficient algorithms that do not use this information.

Backbones. A backbone, for a given communication graph G, is a subnetwork H which forms a connected dominating set of G with asymptotically the same diameter as D and a constant number of internal neighbors in H for every node. Additionally, each node is equipped with specific protocols that allow efficient use of the backbone. More precisely, given a network with a communication graph G, a *backbone* of G is a subnetwork H with the following properties:

1. The nodes of H form a connected dominating set of G and have constant degree.
2. The number of nodes in H is $\mathcal{O}(m)$, where m is the size of the smallest connected dominating set of G.
3. Each node from $G \setminus H$ is associated with exactly one of its neighbors that belongs to H, which is referred to as its *leader*.

There are two protocols $\textsc{Inter}_H$ and $\textsc{Intra}_H$ associated with a network G and its backbone H, that implement inter-cluster communication (among the leaders) and intra-cluster communication (between nodes and their leaders), respectively. These protocols have w the following characteristics:

Protocol $\textsc{Inter}_H$: It is a protocol that makes it possible to exchange messages between each pair of neighbors in H in one multi-round consisting of a_1 rounds. Each message received by a

node $v \in H$ is successfully broadcast to all nodes from G associated with v, for which v is the leader.

Protocol $\textsc{Intra}_H$: It is a protocol which facilitates transmitting a message of each node to its leader in a_2 rounds.

Thes notions and definitions follow [11]. Backbone can be used as a generic tool for many communication and computation task in the network, see [11, 26].

3 Algorithmic tools

We review concepts and tools that are used as components of algorithms.

Grids and boxes. The notions of grids and boxes occur naturally when designing algorithms when nodes are points in Euclidean space. For a constant $b > 0$, we consider a grid of lines parallel to the coordinate axes that partition the space into $b \times b$ boxes, and such that $(0, 0)$ is a point on the grid. More precisely, a *box* includes its left side without the top point and its bottom side without the rightmost point, but not the other two sides. When $b = r/\sqrt{2}$ then the grid is called *pivotal*. This grid has the property that any two nodes in a box are in each other's ranges.

Strongly-selective families. A family $S = (S_0, \ldots, S_{s-1})$ of subsets of $[N]$ is said to be of *length* s. Such an S is a (N, x)-*strongly-selective family*, or (N, x)-*ssf* for short, if, for any non-empty subset Z of $[N]$ such that $|Z| \leq x$ and for any element $z \in Z$, there is a set S_i in S for which $S_i \cap Z = \{z\}$. It is known from [3] that there exists an (N, x)-ssf of length $\mathcal{O}(x^2 \log N)$ for every $x \leq N$. We identify a family of sets $S = (S_0, \ldots, S_{s-1})$ with a broadcast schedule S' such that the i^{th} bit of $S'(v)$ is equal to 1 iff $v \in S_i$. Such a broadcast is said to be performed by *executing* S.

A wireless network based on the SINR model of weak sensitivity could be intuitively interpreted as represented by a unit disc graph with transmission range $1-\varepsilon$ further restricted by a property that only a "dominating" signal from a station within the transmission range could be heard. It follows that efficient algorithms could employ techniques used in unit disc graphs with some transmission schemes that make sets of transmitters relatively "sparse" throughout the network. For the latter we will be using strongly selective families with suitably chosen parameters.

Next, we discuss properties of (N, c)-ssf families that are relevant to the SINR model, where a parameter c will be specified later. The following Proposition 1 is the main technical insight we will use to reason about properties of SINR wireless communication. It will be combined with combinatorial and graph-related methods when analyzing algorithms presented in this work.

Proposition 1 *If there are at most b active nodes per box and they all execute an (N, c)-ssf family, where $c = \Theta(b^3(\varepsilon\mathcal{N})^{-2}) = \Theta(b^3)$, then for each active node there is a round in which the node successfully informs all its neighbors.*

Proof: Suppose that in each box there are at most b active nodes, for some parameter b to be settled later. Let $\eta = \eta(b, \varepsilon\mathcal{N})$ be some suitably large function linear in b and $(\varepsilon\mathcal{N})^{-1}$, which is constant with respect to the transmission graph parameters and the size of the network.

Let us assume that all active nodes transmit according to a (N, c)-ssf, for some sufficiently large parameter $c = \Theta(b\eta^2) = \Theta(b^3(\varepsilon\mathcal{N})^{-2})$. Consider an active node in the network. We observe that a protocol based on an (N, c)-ssf guarantees the existence of a round in which this active node transmits, while no other active node in the surrounding boxes within the Manhattan radius η of the considered node transmits. Indeed, there are $\mathcal{O}(b\eta^2) = \mathcal{O}(b^3(\varepsilon\mathcal{N})^{-2})$ such nodes, and therefore the $(N, \Theta(b^3(\varepsilon\mathcal{N})^{-2}))$-ssf family guarantees that the considered active node has its unique transmission round with respect to the other $\mathcal{O}(b\eta^2)$ active nodes in the respective surrounding boxes within the Manhattan distance η.

Additionally, the total interference from active nodes in the boxes outside the Manhattan distance η from the considered active node and all its neighbors could be made arbitrarily small by selecting a sufficiently large constant multiplied by b in the definition of η. (A similar approach was applied in [17]). Here we also rely on the property that there are at most b active nodes in each box. More precisely, the total interference is bounded from above by

$$\sum_{j\geq\eta} 4j \cdot b \cdot j^{-\alpha} = \mathcal{O}\Big(b\eta^{1-\alpha} \cdot \sum_{j\geq 1} j^{1-\alpha}\Big) = \mathcal{O}\big(b\eta^{1-\alpha}\big) \ .$$

This is because there are at most $4j$ boxes of Manhattan distance $j+1$ from the considered box, each of them with at most b transmitters among active nodes, as only those can transmit, and the signal strength is $\Theta(j^{-\alpha})$ at the considered node from each of them. In the last inequality above, we used the fact that

$$\sum_{j\geq 1} j^{1-\alpha} = \mathcal{O}(1) \ ,$$

for $\alpha > 2$. Observe that $\mathcal{O}\left(b\eta^{1-\alpha}\right)$ could be made smaller than $\varepsilon\mathcal{N}$ by choosing an appropriate function η, following its definition.

An active node informs all its neighbors successfully in the round under consideration, because the signal of the active neighbor is at least $(1+\varepsilon)\beta\mathcal{N}$, by the definition of the hearing range, while the total interference plus noise is at most $\mathcal{N} + \varepsilon\mathcal{N}$. Hence the value of SINR is at least β, which follows from Equation (1). □

In our algorithms, we will use (N, c)-ssf for sufficiently large c, with c being either $c = O(1)$ or $c = \Theta(\log^3 N)$, corresponding to $b = \mathcal{O}(1)$ and $b = \mathcal{O}(\log N)$, and of lengths $\mathcal{O}(\log N)$ and $\mathcal{O}(\log^7 N)$, respectively.

Selectors. Let N, x and y be positive integers such that $y \leq x \leq N$. Let S be a family of subsets of $[N]$. Following [5], we say that S is a (N, x, y)-*selector* if for each set $A \subseteq [N]$ of size $|A| = x$, there are at least y elements in A that can be selected from A by subsets belonging to S. It is also known that for $y = d \cdot x$, where $d \in (0, 1)$ is a constant, there is an (N, x, y)-selector of size $\mathcal{O}(x \log N)$.

Odds and ends. Our algorithms' bounds hold with high probability, which means a probability that is at least $1 - \frac{1}{n^d}$, for some sufficiently large constant $d \geq 1$. For the simplicity of exposition, we assume that the numbers n, N and Δ are all powers of 2; in the general case, these parameters can be rounded up to the nearest power of 2.

4 Broadcast to coordinate start

In this section, we present a randomized algorithm for single-source broadcast when nodes other than the source are allowed to join the execution of an algorithm only after they have received a message. This algorithm is necessary in the uncoordinated-start setting, since all nodes have to be activated. A synchronization of time could be additionally assured by a broadcast protocol, e.g., by updating and forwarding round counter with every message sent by the broadcast protocol. Once such a broadcast is completed, algorithms designed for synchronized start setting can be applied, like the backbone algorithm given in Section 5. We execute a depth-first search on the network by passing a token. Since any (intermediate) node passing the token does not know its immediate neighborhood, this information needs to be learnt. The process of learning neighborhood should be such that at most polylog n rounds may be spent on discovery of a new node. The total time spent may be averaged out and the analysis of our algorithm will be amortized. An aspect playing an important role in the design of our algorithm is that we would like to be parsimonious in use of random bits and supply each node with a number of bits that is poly-logarithmic in n.

We introduce some terms for formal presentation of the protocol. For a node s, a node v is called *viewed-by-s* if the first time v receives a message from some node then it is from s and v has successfully transmitted its label to s; otherwise the node is called *unexplored*. When the token is passed to a node during the search process, it is assigned color *grey*. Once the node does not have any more neighbors to explore, it is assigned color *black* and it returns the token back to the parent from which it received the token for the first time. The search process uses the three routines with the following goals:

1) Estimate the size of the unexplored neighborhood:

   We present a randomized procedure that approximately computes this value in $\mathcal{O}(\log^2 n)$ rounds with high probability.

2) Learn the unexplored neighborhood:

   Given an upper bound on the size of the unexplored neighborhood, this procedure learns the entire unexplored neighborhood of a node.

3) Explore and pass the token:

   This routine executes depth first search on the network. When a token is passed to an unexplored node, it checks whether it has neighboring nodes to which token has not been passed before or whether the token should be returned because no such a neighbor exists. For executing this approach, a node needs first to learn the size of its unexplored neighborhood.

4.1 Estimating the size of an unexplored neighborhood

A fundamental limitation of the SINR model is that it does not allow the Echo() type process, as presented in [19], which has been widely exploited in ad-hoc radio networks. We devise a randomized procedure to approximately estimate the size of the unexplored neighborhood of a node. This procedure is initiated and coordinated by a node s, which, intuitively, is the one with the token in the main algorithm, where only nodes in its neighborhood determined by distances at most $1-\varepsilon$ can participate. The node s tries to estimate the size of the neighborhood by measuring the number of transmissions and comparing it with some threshold value.

Algorithm: Estimate the Size of the Unexplored Neighborhood (ESUN)

Stage 1. The unexplored neighbors of the node execute $\log N$ sub-stages. In sub-stage i, the unexplored neighbors repeatedly transmit with probability $1/2^i$ each, for $d \log N$ rounds, for an appropriately chosen constant d.

Stage 2. Let k be the highest sub-stage in the previous stage for which it is true that the number of successful transmissions, denoted n_k, satisfies $n_k \geq \frac{d \log N}{8}(1-\epsilon)$, for an appropriately chosen constant ϵ. The node s declares an upper bound 2^k on the size of the unexplored neighborhood.

The initiating node s is able to learn its entire unexplored neighborhood in $\mathcal{O}(2^{i+1} \log N)$ rounds, given an upper bound estimate 2^{i+1} with high probability, as we show next:

Lemma 1 *Let the size of the unexplored neighborhood of s be denoted by Δ. Then the upper bound x on the size of the unexplored neighborhood of s declared by s is accurate, that is, the inequalities $\Delta \leq x \leq 16\Delta$ hold with high probability.*

Proof: The parameter Δ can fall in to the following buckets of geometrically increasing sizes

$$[0,1], [2,3], [4,7], \ldots, [2^{i-1}, 2^i - 1], \ldots, [2^{\log N - 1}, 2^{\log N}] \ .$$

Let $\Delta \in [2^{k-1}, 2^k - 1]$. The expected number of transmissions from neighboring nodes μ_k in an execution of the k^{th} sub-stage of Stage 2 of the algorithm is at least $\Delta \cdot \frac{1}{2^k} = \frac{\Delta}{2^k}$. Therefore, we have that $\mu_k \in [1/2, 1)$. Define $X_k^i = 1$ iff a single transmission happens in the i^{th} trial of the k^{th} sub-stage of the algorithm. Then we have that $X_k^i = 1$ when one node transmits and the rest remain silent. The probability that m^{th} node transmits while the rest remain quiet, can be bounded as follows:

$$\frac{1}{2^k}\Big(1 - \frac{1}{2^k}\Big)^{\Delta - 1} \geq \frac{1}{2^k} \cdot \Big(1 - \frac{1}{2^k}\Big)^{2^k \frac{\Delta}{2^k}} \geq \frac{1}{2^k} \cdot \Big(1 - \frac{1}{2^k}\Big)^{2^k} \geq \frac{1}{2^k} \cdot \frac{1}{4} \ .$$

Thus, the probability that one single node transmits while the rest remain silent is at least

$$\Delta \cdot \frac{1}{2^k} \cdot \frac{1}{4} \geq \frac{1}{8} \ .$$

Thus, the following inequality holds: $\Pr[X_k^i = 1] \geq \frac{1}{8}$.

Let the number of successful transmissions received by node s in i^{th} sub-stage of algorithm ESUN be denoted by n_i. If we define

$$X_k = \sum_{i \in [d \log N]} X_k^i \ ,$$

then the following inequality holds:

$$E[X_k] \geq \frac{1}{8} \cdot d \log N \ .$$

Thus, node s expects to hear at least $\frac{d \log N}{8}$ transmissions in k^{th} sub-stage. By a Chernoff bound, we deduce the following fact: the inequality $n_k \geq \frac{(1-\epsilon) d \log N}{8}$ holds with high probability. It follows that the upper bound x declared by the algorithm is at least $2^k \geq \Delta$ with high probability.

It remains to show that this declared upper bound x is not bigger than $2^{k+3} \leq 16\Delta$ with high probability. To this end, consider an integer $k+4 \leq i \leq \log n$ and a sub-stage i in Stage 1 of the algorithm. In a single round of this sub-stage, the probability that there is at least one transmitting station is at most $\Delta \cdot \frac{1}{2^i} \leq \frac{1}{16}$. Thus the expected number of rounds with at least one transmission in sub-stage i is at most $\frac{d \log N}{16}$. Using a Chernoff-bound argument, we deduce that $n_i \leq \frac{(1-\epsilon)d \log N}{8}$ with high probability, for all $i \geq k+4$. Hence the highest sub-stage in which the threshold could be exceeded is $k+3$ with high probability, and thus the inequalities $x \leq 2^{k+3} \leq 16\Delta$ hold. $\square$

4.2 Learning an unexplored neighborhood

Let us assume that we are given an estimated upper bound x on the size Δ of the unexplored neighborhood of a node s such that $\Delta \leq x \leq 16\Delta$. We present a routine for s to learn the entire neighborhood in $\mathcal{O}(x \log N)$ rounds, given this estimate.

Algorithm: Learn Unexplored Neighborhood (LUN)

Let it be given that the number of unexplored neighbors of the initiating node, which have previously not received any message, is at most x with high probability. Let us assign $c = 0.99$. Repeat the two stages below for $j = 0, 1, \ldots, \log_c x$:

Stage 1: Node s asks its unexplored neighbors to execute a $(N, (1-c)^j x, c(1-c)^j \cdot x)$-selector of length $\mathcal{O}((1-c)^j x \log N)$, during which it receives labels of a large fraction of them.

Stage 2: Node s announces the neighbors it received in the previous stage, transmitting them one after another. These nodes will not participate in the execution from now on.

When algorithm LUN is executed, a fraction of the neighborhood of s is learnt in every execution of the two stages, by the definition of a selector, and thus the size of the remaining unlearned neighborhood decreases geometrically. Hence, the lengths of the stages sum up to $\mathcal{O}(x \log N)$.

Similarly as in case of procedure ESUN, procedure LUN does not use any geometric properties of the network, in particular, it does not rely on properties involved in the definition of SINR.

4.3 A token traversal

We use an algorithm, described next, that has a token traverse the network according to the depth-first traversal principles. We refer to colors of nodes used to implement such a traversal. Nodes do not know their neighbors initially. A neighboring node is "viewed" once by a given node to record it a neighbor. All the needed communication is performed by either by referring to Algorithms ESUN and LUN or by one-hope broadcast or unicast. Therefore, this does not rely on geometric nature of the network and in particular on any SINR threshold.

Next we describe a specific implementation of depth-first search.

Algorithm Wireless_DFS

Initially the source possesses the token. The following is executed when a node s receives a token:

Stage 1: Node records the label of the node from which it received the token the first time and assigns itself the color grey.

Stage 2: If the node s has not previously executed the algorithm to estimate the size of its unexplored neighborhood, then s executes Algorithm ESUN to obtain an upper bound on the size of the unexplored neighborhood of s. This estimate is used by s to execute Algorithm LUN. Let R denote the list of unexplored neighbors discovered by s during this execution.

(a) $R \neq \emptyset$: Node s chooses a node $s_1 \in R$ arbitrarily and passes the token to it.

(b) $R = \emptyset$: Node s assigns itself the color black and returns the token to its parent node, from which it received the token for the first time.

Stage 3: Otherwise, s chooses a node from the subset of the neighborhood viewed-by-s, of which the token has not explored yet and passes the token to it. If all the neighbors have been explored by the token, then s colors itself black and returns the token back to the parent node from which it received the token for the first time.

Theorem 1 *Algorithm* WIRELESS_DFS *accomplishes a single-source broadcast from uncoordinated start in* $\mathcal{O}(n \log^2 N)$ *rounds with high probability. The total number of random bits used per node of the network is* $\mathcal{O}(\log^3 N)$.

Proof: The proof of correctness follows from the following facts: (1) Lemma 1 which states that an upper bound on the size of the unexplored neighborhood of a node is correctly estimated with high probability, (2) The correctness of Algorithm LUN for learning the unexplored neighborhood of a node given an upper bound on its size with high probability, and (3) The correctness of the depth first search routine, in which every node of the network is viewed exactly once and connected to this grey node which viewed it.

We only need to prove the bound on the round complexity of the algorithm and the number of random bits used per node of the network. Let us associate the cost of estimating the size of a neighborhood to the node initiating the process and the cost of learning the neighborhood to the individual nodes which were viewed in the process. Each node executes the algorithm for estimating its unexplored neighborhood at most once, spending $\mathcal{O}(\log^2 N)$ rounds in the process, and once a node is viewed by a node it is not viewed by another node. Learning a neighborhood of x nodes takes $\mathcal{O}(x \log n)$ rounds, hence the amortized cost of learning a neighboring node is $\mathcal{O}(\log n)$.

The search routine discovers and explores a node only via the node which viewed it and other edges in the network remain untouched. Thus, the number of edges explored by the search routine is $\mathcal{O}(|V|)$. Lastly, we account the cost of passing the token to a node and of returning it back to the parent to the node itself. Further noting that the token can be passed to a node only once by the node which views it, we have that the round complexity of the algorithm is $\mathcal{O}(n \log^2 N)$ rounds. Note that every node participates in exactly one execution of procedure ESUN, in which it uses $\mathcal{O}(\log^3 N)$ random bits. □

Theorem 1 refers to the number of nodes n in the performance bound, but this parameter n is not known to the nodes, and only N is.

5 Building a backbone from a synchronized start

In this section we develop a randomized algorithm for building a backbone network locally for the setting of synchronized start in which nodes are awake at the start of the protocol. Our main

protocol has a complexity of $\mathcal{O}(\Delta \log N + \log^c N)$ rounds, where c is a positive constant. The costs of protocols INTER_H and INTRA_H associated with building a backbone are $a_1 = \mathcal{O}(\log N)$ and $a_2 = \mathcal{O}(\Delta \log N)$, respectively.

The algorithm consists of two procedures. One of them elects a maximal independent set (MIS) of nodes in the communication graph, which is maximal in the sense of inclusion. The other one elects *connectors*, which are nodes that connect the nodes in a MIS to create the backbone, and computes transmission schedules for protocols INTER_H and INTRA_H associated with the backbone. Next we present these two procedures.

5.1 A procedure for finding maximal independent sets

Due to the assumed model of SINR, the approach used in [4] to find a MIS in the strong connectivity model is not immediately applicable, and we develop a specialized algorithm to accomplish this.

Our procedure to find a MIS is called MIS_SWD. Members of a MIS the procedure identifies are called leaders. At the beginning of the procedure, all nodes have a *worker* status. A status may change in the course of the algorithm to become either a *candidate*, *leader* or *slave*. The last two states are terminal statuses, in that once one of them has been reached, the node keeps this status forever.

The procedure works in phases, parameterized by integer i where $0 \leq i \leq \log \Delta$. The phase for parameter i consists of $\Theta(\log N)$ sub-phases, where the constant hidden in the Theta notation will follow from the analysis. Sub-phase j of phase i consists of two stages described below.

Stage 1: Election of new local leaders. At the start of this stage, each node with status *worker* changes its status to *candidate*, with probability $2^i/\Delta$. The candidates then transmit their names using $(N, \Theta(\log^3 N))$-ssf of length $\mathcal{O}(\log^7 N)$. All nodes record the names from the messages received from candidates and leaders. Candidates that hear a message from at least one other candidate change their status back to worker at the end of the stage, otherwise they change their status to leader.

Stage 2: Enslaving workers by new leaders. All nodes that become leaders in the preceding Stage 1 transmit their own names using an $(N, \Theta(\log^3 N))$-ssf of length $\mathcal{O}(\log^7 N)$. All nodes hearing messages from them record their presence. At the end of the stage, each worker that hears at least one node's name during this stage chooses the leader with the smallest name heard as its *master* node. Simultaneously, the node changes its status to *slave*.

Next we analyze procedure MIS_SWD. By the specification, Stages 1 and 2 take $\mathcal{O}(\log^7 N)$ rounds each. The analysis of correctness relies on the following invariant, which we will prove to hold in the beginning of each phase i with a high probability:

(a) each box of the pivotal grid has at most $\Delta/2^i$ workers in the communication graph, and

(b) each node who has a leader as its neighbor has status slave, and

(c) each node has at most 25 neighbors of leader status and no node has candidate status.

More specifically, we want to show that the invariant holds for every phase i, with a probability at least $1 - (i-1)/n^2$. The invariant for $i = 1$ holds, with probability 1, because of the definition

of Δ and the initial setup with all nodes being workers. Suppose the invariant holds for some $0 \leq i < \log \Delta$ with a probability that is at least $1 - i/n^2$. We show that it also holds for $i+1$ with a probability that is at least $1 - (i+1)/n^2$. To this end, consider a phase i. We start by stating the following useful property concerning the evolution of nodes' status.

Fact 1 *A node can be a candidate only inside Stage 1 of a sub-phase, and therefore the sets of workers in the beginning of sub-phases are monotonically non-increasing in the sense of inclusion.*

Proof: A node can become a candidate only in the beginning of Stage 1. Each candidate can graduate to become either a worker or a leader at the end of Stage 1, and the conditions for these changes are mutually complementary; therefore no candidate remains at the end of Stage 1. The other two states of leader and slave are final, so the claim holds. □

The next lemma proves that the contention in each box, and thus also the total interference, are bounded in each sub-phase.

Lemma 2 *Consider a sub-phase j of phase $i+1$. In the beginning of the sub-phase there are at most $\Delta/2^i$ workers in every box of the pivotal grid with probability at least $1-(i-1)/n^2$. Moreover, $\mathcal{O}(\log N)$ nodes become candidates with the conditional probability of at least $1-1/n^5$.*

Proof: The proof is by induction on the sub-phase numbers j. In the beginning, the first part of the Lemma follows directly from the invariant (a) for phase i. By the first part and an estimate of the probability of becoming a candidate, the expected number of candidates in a box is at most 1. By the Chernoff bound, deviating by a factor $\Theta(\log N)$ from the expectation occurs with a probability that is exponentially small in $\log N$, which means that the $\mathcal{O}(\log N)$ bound holds with high probability. More precisely, we can make it to be $1 - 1/n^5$ by setting a sufficiently large constant in front of $\log N$.

Now we prove the inductive step. Assume that the statement holds in the beginning of a sub-phase $j \geq 1$. During this sub-phase, the number of workers cannot increase, as being a candidate is transient by Fact 1, while other states, namely *leader* and *slave*, are definitive. Therefore, the first part of the Lemma holds at the end of sub-phase j, that is, in the beginning of sub-phase $j+1$ or phase $i+1$, if j is the last sub-phase. The argument for the second claim of the lemma is exactly the same as for sub-phase 1, and follows from a Chernoff bound. □

Based on the Lemma 2 and the stated above properties of an $(N, \Theta(\log^3 N))$-ssf used in the procedure MIS_SWD, we can prove recursively the invariant (a), and later the invariants (b) and (c).

Lemma 3 *If the invariant for i holds then the invariant (a) holds for $i+1$ with the conditional probability at least $1 - 1/n^3$.*

Proof: Assume that the invariant for i holds, in particular, all boxes have no more than $\Delta/2^i$ workers in the beginning of a phase i. Consider such a phase i. Recall that the sets of workers are monotonically non-increasing, by Fact 1. In what follows, we calculate probabilities conditional on that the invariant for i does hold.

We first prove the invariant (a). Consider a box which has more than $\Delta/2^{i+1}$ workers in the beginning of phase i. First observe that during each sub-phase of phase i the conditional probability

that exactly one worker in the box and no other worker in the other 24 closest surrounding boxes, becomes a candidate is a positive constant. Therefore, the probability that there is a sub-phase j in phase i satisfying this condition can be made at least $1 - 1/n^5$, by choosing a sufficiently large constant in front of $\log N$ in the formula on the number of sub-phases in a phase. We consider the following event in this sub-phase j : there is exactly one candidate in the box *and* there is no other candidate in the 24 closest surrounding boxes *and* in every other box there are $\mathcal{O}(\log N)$ candidates. By Lemma 2 and the union bound for the events in boxes, this event holds with a conditional probability at least $1 - 1/n^4$. Hence, conditional on this event, the only candidate in the box does not hear any name during Stage 1, and thus becomes a leader. Moreover, in the subsequent Stage 2, the applied $(N, \Theta(\log^3 N))$-ssf family of length $\mathcal{O}(\log^7 N)$ guarantees the existence of a round in which only the new leader in the box transmits and no other candidate in the surrounding boxes within radius $\Theta(\log N)$ of the considered box transmits. By Proposition 1, the only transmitter from the box successfully informs all its neighbors, in particular all workers in the box who in turn become slaves at the end of Stage 2. To summarize, the conditional probability that there is a sub-phase in which a new leader is elected in the box and it succeeds to enslave all remaining workers in the box is at least $1-1/n^4$. By taking the union of events over all boxes (that have more than $\Delta/2^{i+1}$ workers in the beginning of phase i), the invariant (a) for $i+1$ holds with a probability that is at least $1 - 1/n^3$. □

Lemma 4 *If the invariant for i holds then the invariants (b) and (c) hold for $i+1$ with the conditional probability at least $1 - O(\log N)/n^3$.*

Proof: We assume that the invariant for i holds and calculate probabilities conditional on this event. In order to prove the invariant (b) for $i+1$, notice first that two leaders who are in each other's ranges are elected in the same sub-phase with a probability that is smaller than $1 - 1/n^3$. This is because, by Lemma 2, the number of candidates in each box is $\mathcal{O}(\log N)$ with a probability at least $1 - 1/n^5$. Therefore, the $(N, \Theta(\log^3 N))$-ssf family of length $\mathcal{O}(\log^7 N)$ applied in Stage 1 guarantees the existence of two rounds in which each of these two nodes transmits alone, respectively, and no other candidate in the surrounding boxes within radius $\Theta(\log N)$ of the considered box transmits. By Proposition 1, the only transmitter from the box successfully informs all its neighbors, and thus the two nodes would hear each other during Stage 2 and would transition back to the status of worker, rather than to become leaders. There are at most n^2 such pairs of nodes. Therefore, no two neighboring nodes become leaders in a single sub-phase with a probability that is at least $1-n^2/n^5 = 1-1/n^3$, by the union bound. Since there are $\mathcal{O}(\log N)$ sub-phases, another application of the union bound proves the invariant (b) for $i+1$ with a conditional probability of at least $1 - O(\log N)/n^3$. When the invariant (b) holds, then the invariant (c) is satisfied as well, because there can be at most one leader in a box, as otherwise there would be two neighboring leaders contradicting the invariant (b), and moreover there are at most 25 surrounding boxes within the range of each node, by properties of the pivotal grid. □

Theorem 2 *Procedure* Mis_SWD *constructs a maximal independent set with a probability that is at least $1 - (\log \Delta)/n^2$ in $\mathcal{O}(\log^8 N \log \Delta)$ rounds.*

Proof: The invariant holds for all phases with a probability at least $1 - (\log \Delta)/n^2$, by the union bound applied to all phases based on events and the bounds on probabilities given in Lemmas 3

and 4. Observe now that the end of the last phase could be interpreted as a beginning of an extra phase $\log \Delta + 1$, the invariant for which could be proved based on the behavior of the procedure in its last phase $\log \Delta$. An argument to show this proceeds along the same lines as the analysis of the previous phases in the proofs of Lemmas 3 and 4. The time complexity follows from the length $\mathcal{O}(\log^7 N)$ of the used strongly selective families, and the partition into $\log \Delta$ phases and $\Theta(\log N)$ sub-phases, respectively. □

5.2 A procedure to find connector nodes

We describe a procedure ConnectBB, which selects connector nodes that connect the nodes in the MIS computed by the first procedure Mis_SWD into the backbone (Stage 1), and to compute transmission schedules for protocols Inter_H and Intra_H associated with the backbone (Stage 2).

Initially, we assume that there is some MIS and each node knows locally whether it belongs or not to this set. A node in such a MIS is called a *leader*, to be consistent with the role of leaders in the previously designed procedure Mis_SWD. Each leader is to become a node of the created backbone, so it keeps the record of all other leaders heard during the algorithm as neighbors in the backbone.

Stage 1: Selecting connectors.

This stage consists of four parts. In the first part of this stage, all non-leader nodes learn the names of their leader neighbors. In order to accomplish this, the leaders transmit their own names using a (N, c)-ssf of length $\mathcal{O}(\log N)$, for a sufficiently large constant $c > 0$, to be determined in the analysis. Every node that is not in the MIS records the first 25 different names heard, or as many were heard if fewer than 25 names were heard. The j-th heard name is called the *j-th leader* of the node.

During the second part, each leader learns about other leaders of distance 3 in the communication graph and choose a constant number of connectors via which it can reach the other leaders. Note that if one connects all pairs of nodes in a MIS that are of distance at most 3 from each other, by adding nodes on arbitrarily selected shortest path between them, then the obtained set of nodes is a connected dominating set. This is done by having each leader (a node in MIS) communicate with its neighbors and nodes of distance two in part 1. This starts with each node v not belonging to MIS choosing a random number t_v in $\Theta(\Delta)$. Let $x = O(\log^7 N)$ be the length of an $(N, \Theta(\log^3 N))$-ssf, which is the same family as in the efficient implementation of procedure Mis_SWD. Proposition 1 is used to argue that the family succeeds to propagate messages from active nodes in boxes to their neighbors, provided there are $\mathcal{O}(\log N)$ active nodes per box. This is organized as two sub-parts, both of a similar structure. During the first sub-part, two adjacent non-MIS nodes exchange information about their neighbors that are leaders. During the second sub-part, each node passes this knowledge along with the information about its own neighbors that are leaders to the neighbors that are leaders, of whom there are at most 25.

We start from specifying the first sub-part.

Let us partition time into consecutive blocks of length x, during which executions of some $(N, \Theta(\log^3 N))$-ssf are done by subsets of nodes. Additionally, we join blocks into consecutive groups of 25 blocks each. Then node v is active during group t_v in every round in $[25xt_v + 1, \ldots, 25x(t_v + 1)]$, in the following sense: in the j-th block of group t_v, where $0 \leq j \leq 24$, (in

rounds $[25xt_v + jx + 1, \ldots, 25xt_v + (j+1)x)$, node v transmits its j-th leader and its own name using an $(N, \Theta(\log^3 N))$-ssf. Each node w records all the names heard in this execution. For every name that u heard, node w recalls the round number f_u when the name of u was heard and the node g_u who sent it.

Next, we describe the second sub-part.

It proceeds similarly as the first sub-part, with the following modifications. First, each group consists of 49 blocks rather than 25 blocks. This is because it is an upper bound on the number of MIS nodes that are two hops away from a node, by geometry of the plane. Second, each block is associated with some known leader, either in one-hop neighborhood, as learned in part 1 of Stage 1, or in two-hop neighborhood, as learned in the first sub-part of part 2 of Stage 1. In the latter case, such leader's name, say u, is transmitted along with its associated forwarding node g_u and the successful round number f_u. Third, at the end of this sub-part, each leader w additionally records all leaders' names u heard in this sub-part, together with one or two node(s) associated with this name, and one or two, respectively, round number(s) when successful transmissions between the node(s) and the leader took place. In case there are many intermediate nodes or pairs thereof, an arbitrary selection of one such a configuration per a leader's name is made: for each leader's name u, let $g_{u,1}$ denote a one-hop neighbor and $g_{u,2}$ be two-hop neighbor associated with this name, and let $f_{u,i}$ be associated with a successful transmission round of node $g_{u,i}$, for $i = 1, 2$. In case there is only one connector associated with a leader of name u, we denote it by g_u and its successful round number by f_u, to simplify notation. We call these nodes *connectors*, as they can forward messages between pairs of leaders.

In the third part of this stage, leaders inform connectors chosen in part 2. Similarly as in part 2, this needs to be done in two sub-parts, as connectors of distance 1 and then of distance 2 need to be informed. In the first sub-part, every leader node w transmits according to a (N, c)-ssf of length $\mathcal{O}(\log N)$, for a sufficiently large constant $c > 0$ to be determined in the analysis, precisely 121 times, one after another. The number 121 is an upper bound on the number of leaders nodes in any 3-hop neighborhood, by a geometric argument. Proposition 1 is used again, for a constant number of active nodes (leaders) per box, to guarantee a successful message exchange between neighboring pairs of leaders. In the j-th execution of a (N, c)-ssf, node w transmits its j-th heard leader u's name along with its associated connector g_u and round value f_u, or a pair of connectors $g_{u,1}, g_{u,2}$ and round numbers $f_{u,1}, f_{u,2}$. Upon receiving such a message containing only one connector g_u and a value f_u, connector g_u records that it is a connector from leader u to leader w, as well as the successful round f_u when forwarding took place. Upon receiving such a message, containing two connectors $g_{u,1}, g_{u,2}$ and values $f_{u,1}, f_{u,2}$, connector $g_{u,1}$ records that it is a connector from leader u, via node $g_{u,2}$, to leader w, as well as the successful round $f_{u,1}$ when the forwarding from it to the leader w took place. In the second sub-part, only such connectors $g_{u,1}$ are active in relaying messages. Similarly as in the first sub-part, they use a (N, c)-ssf of length $\mathcal{O}(\log N)$, for a sufficiently large constant $c > 0$ to be determined in the analysis, precisely 121 times one after another, and in the j-th execution of the (N, c)-ssf, node $g_{u,1}$ transmits its j-th heard associated connector $g_{u,2}$ for the leader u, along with its associated pair of leaders u, w and round value $f_{u,2}$. Upon receiving such a message, connector $g_{u,2}$ records that it is a connector from the leader u to the leader w, via node $g_{u,1}$, as well as the successful round $f_{u,2}$ when the forwarding from it to the connector $g_{u,1}$ took place. A node becomes a part of the backbone when it is a connector to at least one pair of leaders.

Finally, in part four all backbone nodes, both leaders and connectors, learn their neighborhoods in the backbone. In order to do it, they transmit their own names using a (N, c)-ssf of length $\mathcal{O}(\log N)$, for a sufficiently large constant $c > 0$ to be determined in analysis. Proposition 1 guarantees a successful message exchange between neighboring pairs of backbone nodes, as their density per box is constant. Each backbone node records all the receives names as its neighbors in the backbone.

The four parts of this stage take time $\mathcal{O}(\Delta \log^7 N)$ in total.

Stage 2: Computing schedules for protocols Inter$_H$ and Intra$_H$.

Protocol Inter$_H$ associated with this backbone structure schedules all leaders and connectors to transmit using some (N, c)-ssf of length $\mathcal{O}(\log N)$, for sufficiently large constant $c > 0$, to be determined in analysis. Similarly, as in the fourth part of Stage 1, Proposition 1 guarantees a successful message exchange between neighboring pairs of backbone nodes, as their density per box is constant. One run is executed to confirm neighboring backbone nodes. Additionally, if a non-backbone node does not have its associated backbone node, which is a master node, then it selects the smallest among the heard names in the backbone as its master node.

Protocol Intra$_H$ is computed as follows. Recall that we refer to the length of an $(N, \Theta(\log^3 N))$-ssf by $x = \mathcal{O}(\log^7 N)$. Each non-backbone node selects a random number t in the range $y = \Theta(\Delta \log N)$, with a sufficiently large constant in front of Δ, and it transmits its own name with the name of its master according to an $(N, \Theta(\log^3 N))$-ssf in rounds $[t \cdot x + 1, t \cdot x + x]$. Intuitively, the random component allows all slaves to avoid collisions at their master with high probability, while the $(N, \Theta(\log^3 N))$-ssf component in the product allows to avoid clashes between a number of $\mathcal{O}(\log N)$ slaves of different masters with high probability, where the density of masters is constant per box. A master that belongs to the backbone records the name of the slave from which it receives such message. After time $y \cdot x$, each backbone node orders all its slave nodes and proceeds as follows. For each slave node on the list, it executes some (N, c)-ssf of length $z = \mathcal{O}(\log N)$, for sufficiently large constant $c > 0$, to be determined by analysis, transmitting the slave node's name whenever scheduled. These executions occur for every node on the list one by one. This part is completed within Δz rounds, as at most Δ slave nodes are on the list. Each slave node v remembers its position σ_v on the list of slave nodes of its master, upon receiving a message addressed to it from the master. Finally, protocol Intra$_H$ is defined as follows: every non-backbone node v transmits according to a (N, c)-ssf of length $z = O(\log N)$ in rounds $[\sigma_v \cdot z + 1, \sigma_v \cdot z + z]$. The component σ_v in the product allows scheduling all slaves to avoid collisions at their joint master, as computed above, while the (N, c)-ssf component in the product allows to avoid clashes between a constant number of slaves of different masters.

Stage 2 takes $\mathcal{O}(\Delta \log^7 N)$ rounds. Note that both Inter$_H$ and Intra$_H$ are deterministic, though the latter was pre-computed by resorting to randomization.

Theorem 3 *The distributed algorithm* ConnectBB *computes a backbone from a given MIS, works in $\mathcal{O}(\Delta \log^7 N)$ rounds with high probability, and the associated protocols have lengths $a_1 = O(\log N)$ and $a_2 = O(\Delta \log N)$.*

Proof: Consider Stage 1. The first part guarantees that each non-leader gets names of all its neighbors in the MIS. There are at most 25 such neighbors, by the definition of MIS in unit disc graphs. Hence, a properly selected constant c in a (N, c)-ssf of length $\mathcal{O}(\log N)$ is sufficient to achieve this task, by Proposition 1.

During the second part of Stage 2, each leader obtains either direct or relayed information from each of its one- or two-hop neighbors, about their at most 121 leaders, each of distance at most 3 from the leader. By letting each non-leader node to choose a random number in range $\Theta(\Delta)$, for a sufficiently large constant in front of Δ, it is assured that in every box of the pivotal grid $\mathcal{O}(\log N)$ nodes are selected the same round, with high probability. In order to allow each of these nodes to transmit, an $(N, \Theta(\log^3 N))$-ssf of length $x = O(\log^7 N)$ can be used, by Proposition 1 for $b = \mathcal{O}(\log N)$ active nodes per box with high probability (this is similar to the lines of the proof of Lemma 3). The approach is used in both sub-parts of this part. Additionally, the factor of 25 guarantees that in the first sub-part such situations will occur at least 25 times, so each non-leader node will be able to successfully transmit each of its neighboring leaders' names. These names will be successfully received by all their neighbors, including leaders and non-leaders. There might be at most 49 of leaders of one or two hop distance away from a node, due to the definition of MIS and metric properties of the 2D Euclidean space. In the second sub-part, the newly learned information about leaders of one-hop and two-hop distance away, along with the name of an arbitrarily chosen relaying node, is passed once again to all neighbors, using the general technique based on ssf and its properties proved in Proposition 1, as described above. In this way, each leader learns at most 121 other leaders of distance at most 3, together with at most two arbitrary chosen connectors serving as relaying nodes.

In part three, each leader needs to send back the information about what it learned about other leaders of distance at most 3 away, to at most two connectors. To do it, the leaders execute an (N, c)-ssf of length $\mathcal{O}(\log N)$ in the first sub-part, similarly as in part one, 121 times, each time for a different such leader's name, from at most 121 names. In the second part, each connector which receives such a message addressed to it in which it occurs with another connector, relays this message to the other connector using the same procedure as in the first sub-part. Here we note again that each node could be chosen as a connector to at most 121 pairs of leaders, hence repeating the (N, c)-ssf 121 times, each time for a different pair of leaders for which it was chosen as a connector, is sufficient for relaying all such messages to all two-hop connectors. As already specified in the description of the algorithm, Proposition 1 is used, for a constant number b of active nodes (leaders, connectors) per box, to guarantee a successful message exchange between the neighboring pairs of active nodes.

This is repeated once more in part four to assure that both the leaders and the connectors will know which other leaders/connectors are in their range. The argument for correctness is similar to the one in part one, except that instead of at most one leader in each box of the pivotal grid we have a constant number of backbone nodes (leaders and connectors) in a box.

By the definition of a maximal independent set, all non-backbone nodes are in the range of some backbone nodes (in particular, some leaders), while every two leaders of distance at most twice bigger than the transmission range are connected by a single connector, which creates a connected dominating set (CDS) of asymptotically optimal diameter. The size is also asymptotically optimal, as in each box there is at most constant number of backbone nodes while for every 25 inhabited boxes there must be at least one node taken to a size-minimum CDS.

Next, we argue about computing protocols $\textsc{Inter}_H$ and $\textsc{Intra}_H$ associated with the backbone. This takes place in Stage 2 of procedure ConnectBB. The schedule for $\textsc{Inter}_H$ is correct based on the same argument as used in part four of Stage 1. Additionally, one "control" execution of $\textsc{Inter}_H$ allows to compute by each non-backbone node its master in the backbone, as it can hear

all its neighbors in the backbone as there is a constant number of them.

Regarding a computation of a schedule for protocol INTRA_H, the argument is a simplified version of the one for part two of Stage 1. It guarantees that a random number selection combined with properly choosing a strongly selective family allows each backbone node to hear all the names of neighboring non-backbone nodes, with high probability. Based on that knowledge, acknowledgements are sent by backbone nodes to their non-backbone neighbors, each using the same (N, c)-ssf as in the schedule of INTER_H. All non-backbone nodes receive acknowledgements. Based on their position on the list, non-backbone nodes will be scheduled to transmit in INTRA_H, also using an (N, c)-ssf, after all previous nodes on the list are done. The argument for correctness resorts to the property that each execution of an (N, c)-ssf is associated with a set of pairs involving backbone and non-backbone nodes. Hence, due to a constant density per boxes of the pivotal grid, applying the strongly selective family guarantees successful deliveries of each of these pairs, see Proposition 1. This means that there is a round when only transmitter in this pair is transmitting in close proximity, while an arbitrarily small interference from the further distance is guaranteed to hold due to a bounded density of transmitters.

The length of the computation follows directly from the lengths of all parts. The lengths of protocols INTER_H and INTRA_H follow directly from their specifications. In particular, from the length $\mathcal{O}(\log N)$ of the used strongly selective family and by the fact that protocol INTRA_H uses it up to Δ times. □

An algorithm which first executes MIS_SWD followed by CONNECTBB is called BACKBONE. It computes a backbone. The correctness follows by combining Theorem 2 with Theorem 3.

Corollary 1 *Algorithm* BACKBONE *finds a backbone in* $\mathcal{O}(\Delta \log^7 N)$ *rounds with a probability that is at least* $1 - 1/n$*, for some positive constant* a*, and the associated protocols have lengths* $a_1 = O(\log N)$ *and* $a_2 = O(\Delta \log N)$*.*

The computed protocol INTER_H allows message propagation among the leaders in $\mathcal{O}(D \log N)$ rounds, and protocol INTRA_H allows bi-directional complete communication between backbone master nodes and their slaves outside the backbone in $\mathcal{O}(\Delta \log N)$ rounds, with high probability.

6 A backbone construction for a partly synchronized start

In this section, we combine the results of the previous two sections in order to obtain more general distributed backbone construction, called EMULATED_DFS_FOR_BACKBONE, in case when only an arbitrary set of nodes is activated in the beginning. Let $S \neq \emptyset$ denote this set of nodes.

First, nodes in S execute MIS_SWD from Section 5. During this execution, if a node not in S receives a message, it simply ignores it and so does not join the MIS_SWD execution. Note that there is at least one MIS node in the whole network, as the initial set S of awoken nodes is non-empty. Second, each node selected to MIS, that is, a leader node, initiates WIRELESS_DFS execution with itself as a source, see Section 4. The goal is to carry on these executions until a token's owner learns about other token of higher source name; in such a case it ceases the execution for the currently owned token, but can join other executions in the future, for the largest token heard so far. With such a modification, and under the assumption that executions of WIRELESS_DFS

for different sources do not interrupt each other, except of ceasing whenever the token owner learns about a token of higher source name, it follows that the execution of Wireless_DFS initiated by the MIS node of a highest name will successfully visit all nodes, and thus assure that they become active, in $\mathcal{O}(n \log^2 N)$ rounds, by Theorem 1. Hence, after this execution, algorithm Backbone from Section 5 could be executed by all nodes, constructing a backbone in $\mathcal{O}(\Delta \log^7 N)$ rounds with high probability, by Corollary 1.

It remains to show how to emulate parallel executions of Wireless_DFS so that all messages and computation, which are done in a single execution, as analyzed in Section 4, could be done in our setting. In order to do it, we need to accomplish the following two things:

First, emulate each original communication round of Wireless_DFS by using a (N, c)-ssf of length $\mathcal{O}(\log N)$, for a sufficiently large constant $c > 0$, as in Stage 1 of procedure ConnectBB. More precisely, each node that chooses to transmit in original execution of Wireless_DFS, transmits according to its transmission pattern determined by an (N, c)-ssf. Using an ssf suffices if only we maintain a constant density of transmitting nodes per box in the modified algorithm. All messages successfully delivered in original executions will also be delivered in the modified one, see also the second part of the modification. In the beginning, only source nodes have their own source names, while others have none. Upon receiving messages during this emulation procedure, a node decides on the channel feedback from the emulated communication round by checking the following sequence of conditions:

(i) if no message received then the feedback is empty;

(ii) otherwise, if there is a message received with a higher source name, or if the node has no source name yet, then the node ceases its execution for the current source name and updates its stored source name to the highest received source name; in such a case, if there were at least two different senders for the newly updated source name, the node becomes inactive;

(iii) if the node is active and there was exactly one sender for the currently stored source name, then the message from this sender is the feedback from the channel; otherwise, if the node is active, the feedback is empty.

Second, in order to make the above emulation successful, we need to keep an invariant that in each neighborhood there is a *constant* number of different source names stored in active neighboring nodes. (Here being successful means that all messages that would be received from transmitting neighbors in the communication in case of a single execution of Wireless_DFS are received during the emulation, and thus a correct feedback could be computed.) This is sufficient for an (N, c)-ssf of length $\mathcal{O}(\log N)$, for a sufficiently large constant $c > 0$, to allow each originally successful transmission in a single source name execution to go through, by Proposition 1.

To conclude, observe that this constant bound on the number of different active source names in any neighborhood follows by induction on the number of emulated original rounds. The argument is based on geometric properties of the plane, namely, that a single neighborhood can be covered by a constant number of other neighborhoods, actually 25, and by the feedback rule (ii) above that a node always chooses the highest heard source name as its own. The base of induction holds by the fact that original source names constitute a MIS.

Theorem 4 *Algorithm* EMULATED_DFS_FOR_BACKBONE *creates a backbone network for* any *non-empty set of initially activated nodes, in* $\mathcal{O}(n \log^2 N + \Delta \log^7 N)$ *rounds with high probability. The total number of random bits per node is poly-logarithmic in* N*, and the associated backbone protocols have lengths* $a_1 = \mathcal{O}(\log N)$ *and* $a_2 = \mathcal{O}(\Delta \log N)$*.*

Proof: We combine Theorem 1, regarding WIRELESS_DFS with the $\mathcal{O}(\log N)$ rounds for deterministic emulation, with Corollary 1 regarding algorithm BACKBONE, and also its part MIS_SWD used in this section to choose a MIS of original source names. □

7 Conclusion

This work extended the results obtained by designing algorithms when nodes can refer to their geographic coordinates, obtained by Jurdziński and Kowalski [11] and Jurdziński et al. [17], to algorithms executed by nodes without such capability, in the model of weak-sensitivity and weak-connectivity of wireless networks. We showed that a similar performance can be achieved for broadcasting and building backbones when nodes do *not* use their coordinates.

We developed a new combinatorial-based technique that, combined with other graph-based and geometric based procedures, allows to overcome the lack of spatial coordinates and other capabilities. The analysis of the developed algorithms is carried out in the two-dimensional Euclidean space only for the sake of clear presentation and justification of algorithmic ideas, but it can be extended to more general metric spaces, such as bounded-growth ones.

In view of the lower bound $\Omega(D\Delta)$ on time performance of broadcast and wake-up in wireless networks with uncoordinated start, shown in [17], the algorithm for building a backbone we developed demonstrates that preprocessing a network (in order to synchronize the nodes, so that they can start an execution simultaneously) decreases the expected-time performance requirements for some communication tasks in wireless networks that simultaneously are of weak sensitivity and weak connectivity.

References

[1] C. Avin, Y. Emek, E. Kantor, Z. Lotker, D. Peleg, and L. Roditty. SINR diagrams: Convexity and its applications in wireless networks. *Journal of the ACM*, 59(4):18, 2012.

[2] M. H. L. Bodlaender and M. M. Halldórsson. Beyond geometry: towards fully realistic wireless models. In *Proceedings of the* 33*rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 347–356, 2014.

[3] A. E. F. Clementi, A. Monti, and R. Silvestri. Distributed broadcast in radio networks of unknown topology. *Theoretical Computer Science*, 302(1-3):337–364, 2003.

[4] S. Daum, S. Gilbert, F. Kuhn, and C. C. Newport. Broadcast in the ad hoc SINR model. In *Proceedings of the* 27*th International Symposium on Distributed Computing (DISC)*, LNCS 8205, pages 358–372, 2013.

[5] A. De Bonis, L. Gąsieniec, and U. Vaccaro. Optimal two-stage algorithms for group testing problems. *SIAM Journal on Computing*, 34(5):1253–1270, 2005.

[6] F. Fuchs and D. Wagner. On local broadcasting schedules and CONGEST algorithms in the SINR model. In *Proceedings of the 9th International Symposium on Algorithms and Experiments for Sensor Systems, Wireless Networks and Distributed Robotics (ALGOSENSORS)*, LNCS 8243, pages 170–184, 2013.

[7] A. J. Goldsmith and S. B. Wicker. Design challenges for energy-constrained ad hoc wireless networks. *IEEE Wireless Communications*, 9(4):8–27, 2002.

[8] O. Goussevskaia, T. Moscibroda, and R. Wattenhofer. Local broadcasting in the physical interference model. In *Proceedings of the DIALM-POMC Joint Workshop on Foundations of Mobile Computing*, pages 35–44, 2008.

[9] O. Goussevskaia, Y. A. Pignolet, and R. Wattenhofer. Efficiency of wireless networks: Approximation algorithms for the physical interference model. *Foundations and Trends in Networking*, 4(3):313–420, 2010.

[10] M. M. Halldórsson and P. Mitra. Towards tight bounds for local broadcasting. In *Proceedings of the 8th ACM International Workshop on Foundations of Mobile Computing (FOMC)*, page 2, 2012.

[11] T. Jurdziński and D. R. Kowalski. Distributed backbone structure for algorithms in the SINR model of wireless networks. In *Proceedings of the 26th International Symposium on Distributed Computing (DISC)*, LNCS 7611, pages 106–120, 2012.

[12] T. Jurdziński and D. R. Kowalski. Distributed randomized broadcasting in wireless networks under the SINR model. In M.-Y. Kao, editor, *Encyclopedia of Algorithms*. Springer US, 2014.

[13] T. Jurdziński, D. R. Kowalski, M. Różański, and G. Stachowiak. Distributed randomized broadcasting in wireless networks under the SINR model. In *Proceedings of the 27th International Symposium on Distributed Computing (DISC)*, LNCS 8205, pages 373–387, 2013.

[14] T. Jurdziński, D. R. Kowalski, M. Różański, and G. Stachowiak. On the impact of geometry on ad hoc communication in wireless networks. In *Proceedings of the 33rd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 357–366, 2014.

[15] T. Jurdziński, D. R. Kowalski, M. Różański, and G. Stachowiak. On setting-up asynchronous ad hoc wireless networks. In *Proceedings of the 2015 IEEE Conference on Computer Communications (INFOCOM)*, pages 2191–2199, 2015.

[16] T. Jurdziński, D. R. Kowalski, and G. Stachowiak. Distributed deterministic broadcasting in uniform-power ad hoc wireless networks. In *Proceedings of the 19th International Symposium on Fundamentals of Computation Theory (FCT)*, LNCS 8070, pages 195–209, 2013.

[17] T. Jurdziński, D. R. Kowalski, and G. Stachowiak. Distributed deterministic broadcasting in wireless networks of weak devices. In *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming (ICALP)*, LNCS 7966, pages 632–644, 2013.

[18] T. Kesselheim and B. Vöcking. Distributed contention resolution in wireless networks. In *Proceedings of the* 24*th International Symposium on Distributed Computing (DISC)*, LNCS 6343, pages 163–178, 2010.

[19] D. R. Kowalski and A. Pelc. Time of deterministic broadcasting in radio networks with local knowledge. *SIAM Journal on Computing*, 33(4):80–891, 2004.

[20] S. P. Reddy, D. R. Kowalski, and S. Vaya. Multi-broadcasting under the SINR model. *CoRR*, abs/1504.01352, 2015.

[21] C. Scheideler, A. W. Richa, and P. Santi. An $O(\log n)$ dominating set protocol for wireless ad-hoc networks under the physical interference model. In *Proceedings of the* 9*th ACM Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 91–100, 2008.

[22] S. Schmid and R. Wattenhofer. Algorithmic models for sensor networks. In *Proceedings of the* 20*th IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2006.

[23] D. Yu, Q. Hua, Y. Wang, H. Tan, and F. C. M. Lau. Distributed multiple-message broadcast in wireless ad-hoc networks under the SINR model. In *Proceedings of the* 19*th International Colloquium on Structural Information and Communication Complexity (SIROCCO)*, LNCS 7355, pages 111–122. Springer, 2012.

[24] D. Yu, Q.-S. Hua, Y. Wang, J. Yu, and F. C. M. Lau. Efficient distributed multiple-message broadcasting in unstructured wireless networks. In *Proceedings of the* 32*nd IEEE International Conference on Computer Communications (INFOCOM)*, pages 2427–2435, 2013.

[25] D. Yu, Y. Wang, Q.-S. Hua, and F. C. M. Lau. Distributed local broadcasting algorithms in the physical interference model. In *Proceedings of the* 7*th IEEE International Conference on Distributed Computing in Sensor Systems (DCOSS)*, pages 1–8, 2011.

[26] J. Yu, N. Wang, G. Wang, and D. Yu. Connected dominating sets in wireless ad hoc and sensor networks - A comprehensive survey. *Computer Communications*, 36(2):121 – 134, 2013.